arXiv:1412.6039v3 [stat.ML] 22 Feb 2015

GENERATIVE DEEP DECONVOLUTIONAL LEARNING

Yunchen Pu, Xin Yuan and Lawrence Carin
Department of Electrical and Computer Engineering, Duke University, Durham, NC, 27708, USA
{yunchen.pu,xin.yuan,lcarin}@duke.edu

ABSTRACT

A generative model is developed for deep (multi-layered) convolutional dictionary learning. A novel probabilistic pooling operation is integrated into the deep model, yielding efficient bottom-up (pre-training) and top-down (refinement) probabilistic learning. After learning the deep convolutional dictionary, testing is implemented via deconvolutional inference. To speed up this inference, a new statistical approach is proposed to project the top-layer dictionary elements to the data level. Following this, only one layer of deconvolution is required during testing. Experimental results demonstrate powerful capabilities of the model to learn multi-layer features from images, and excellent classification results are obtained on the MNIST and Caltech 101 datasets.

1 INTRODUCTION

Convolutional networks, introduced in Lecun et al. (1998), have demonstrated excellent performance on image classification and other tasks. There are at least two key components of this model: computational efficiency manifested by leveraging the convolution operator, and a deep architecture, in which the features of a given layer serve as the inputs to the next layer above. Since that seminal contribution, much work has been undertaken on improving deep convolutional networks (Lecun et al., 1998), deep deconvolutional networks (Zeiler et al., 2010), convolutional deep restricted Boltzmann machines (Lee et al., 2009), and on Bayesian convolutional dictionary learning (Chen et al., 2013), among others.

An important technique employed in these deep models is *pooling*, in which a contiguous block of features from the layer below are mapped to a single input feature for the layer above. The pooling step manifests robustness, by minimizing the effects of variations due to small shifts, and it has the advantage of reducing the number of features as one moves higher in the hierarchical representation (possibly mitigating over-fitting). Methods that have been considered include average and maximum pooling, in which the single feature mapped as input to the layer above is respectively the average or maximum of the corresponding block of features below. Average pooling may introduce blur to learned filters (Zeiler & Fergus, 2013), and use of the maximum ("max pooling") is widely employed. Note that average and max pooling are deterministic. Stochastic pooling proposed by Zeiler & Fergus (2013) and the probabilistic max-pooling used by Lee et al. (2009) often improve the pooling process. The use of stochastic pooling is also attractive in the context of developing a generative model for the deep convolutional representation, as highlighted in this paper. Specifically, we develop a deep generative statistical model, which starts at the highest-level features, and maps these through a sequence of layers, until ultimately mapping to the data plane (e.g., an image). The feature at a given layer is mapped via a multinomial distribution to one feature in a block of features at the layer below (and all other features in the block at the next layer are set to zero). This is analogous to the method in Lee et al. (2009), in the sense of imposing that there is *at most* one non-zero activation within a pooling block. As we demonstrate, this yields a generative statistical model with which Bayesian inference may be readily implemented, with all layers analyzed jointly to fit the data.

We use bottom-up pretraining, in which initially we sequentially learn parameters of each layer one at a time, from bottom to top, based on the features at the layer below. However, in the refinement phase, all model parameters are learned jointly, top-down. Each consecutive layer in the model is locally conjugate in a statistical sense, so learning model parameters may be readily performed using sampling or variational methods. We here develop a Gibbs sampler for learning, with the goal of obtaining a maximum *a posterior* (MAP) estimate of the model parameters, as in the original paper on Gibbs sampling (Geman & Geman, 1984) (we have found it unnecessary, and too expensive, to attempt an accurate estimate of the full posterior). The Gibbs sampler employed for parameter learning may be viewed

as an alternative to typical optimization-based learning (Lecun et al., 1998; Zeiler et al., 2010), making convenient use of the developed generative statistical model.

The work in Zeiler et al. (2010); Chen et al. (2013) involves learning convolutional dictionaries, and at the testing phase one must perform a (generally) expensive nonlinear deconvolution step at each layer. In Kavukcuoglu et al. (2010) convolutional dictionaries are also learned at the training stage, but one also simultaneously learns a convolutional filterbank and nonlinear function. The convolutional filterbank can be implemented quickly at test (no nonlinear deconvolutional inversion) and, linked with the nonlinear function, this computationally efficient testing step is meant to approximate the devonvolutional network.

We propose an alternative approach to yield fast inversion at test, while still retaining an aspect of the nonlinear deconvolution operation. As detailed below, in the learning phase, we infer a deep hierarchy of convolutional dictionary elements, which if handled like in Zeiler et al. (2010), requires joint deconvolution at each layer when testing. However, leveraging our generative statistical model, the dictionary elements at the top of the hierarchy can be mapped through a sequence of linear operations to the image/data plane. At test, we only employ the features from the top layer in the hierarchy, mapped to the data plane, and therefore only a single layer of deconvolution need be applied. This implies that the test-time computational cost is independent of the number of layers employed during the learning phase.

This paper makes three contributions: (i) rather than employing beta-Bernoulli sparsity at each layer of the model separately, as in Chen et al. (2011; 2013), the sparsity is manifested via a multinomial process between layers, constituting stochastic pooling, and allowing coupling all layers of the deep model when learning; (ii) the stochastic pooling manifests a proper top-down generative model, allowing a new means of mapping high-level features to the data plane; and (iii) a novel form of testing is employed with deep models, with the top-layer features mapped to the data plane, and deconvolution only applied once, directly with the data. This methodology yields excellent performance on image-recognition tasks, as demonstrated in the experiments.

2 Modeling Framework

The proposed model is applicable to general data for which a convolutional dictionary representation is appropriate. One may, for example, apply the model to one-dimensional signals such as audio, or to two-dimensional imagery. In this paper we focus on imagery, and hence assume two-dimensional signals and convolutions. Gray-scale images are considered for simplicity, with straightforward extension to color.

2.1 Single-Layer Convolutional Dictionary Learning

Assume N gray-scale images $\{\mathbf{X}^{(n)}\}_{n=1,N}$, with $\mathbf{X}^{(n)} \in \mathbb{R}^{N_x \times N_y}$; the images are analyzed jointly to learn the convolutional dictionary $\{\mathbf{D}^{(k)}\}_{k=1,K}$. Specifically consider the model

$$\mathbf{X}^{(n)} = \sum_{k=1}^{K} \mathbf{D}^{(k)} * (\mathbf{Z}^{(n,k)} \odot \mathbf{W}^{(n,k)}) + \mathbf{E}^{(n)}, \tag{1}$$

where $*$ is the convolution operator, $\odot$ denotes the Hadamard (element-wise) product, the elements of $\mathbf{Z}^{(n,k)}$ are in $\{0,1\}$, the elements of $\mathbf{W}^{(n,k)}$ are real, and $\mathbf{E}^{(n)}$ represents the residual. $\mathbf{Z}^{(n,k)}$ indicates which shifted version of $\mathbf{D}^{(k)}$ is used to represent $\mathbf{X}^{(n)}$. Considering $\mathbf{D}^{(k)} \in \mathbb{R}^{n_{d_x} \times n_{d_y}}$ (typically $n_{d_x} \ll N_x$ and $n_{d_y} \ll N_y$), the corresponding weights $\mathbf{Z}^{(n,k)} \odot \mathbf{W}^{(n,k)}$ are of size $(N_x - n_{d_x} + 1) \times (N_y - n_{d_y} + 1)$.

Let $w_{i,j}^{(n,k)}$ and $z_{i,j}^{(n,k)}$ represent elements (i,j) of $\mathbf{Z}^{(n,k)}$ and $\mathbf{W}^{(n,k)}$, respectively. Within a Bayesian construction, the priors for the model may be represented as (Paisley & Carin, 2009):

$$z_{i,j}^{(n,k)} \sim \text{Bernoulli}(\pi_{i,j}^{(n,k)}), \qquad \pi_{i,j}^{(n,k)} \sim \text{Beta}(a_0, b_0), \tag{2}$$

$$w_{i,j}^{(n,k)} \sim \mathcal{N}(0, \gamma_w^{-1}), \quad \mathbf{D}^{(k)} \sim \mathcal{N}(0, \gamma_d^{-1}\mathbf{I}), \quad \mathbf{E}^{(n)} \sim \mathcal{N}(0, \gamma_e^{-1}\mathbf{I}), \tag{3}$$

$$\gamma_w \sim \text{Ga}(a_w, b_w), \quad \gamma_d \sim \text{Ga}(a_d, b_d), \quad \gamma_e \sim \text{Ga}(a_e, b_e), \tag{4}$$

where $i = 1, \ldots, N_x - n_{d_x} + 1$; $j = 1, \ldots, N_y - n_{d_y} + 1$, $\text{Ga}(\cdot)$ denotes the gamma distribution, $\mathbf{I}$ represents the identity matrix, and $\{a_0, b_0, a_w, b_w, a_d, b_d, a_e, b_e\}$ are hyperparameters, for which default settings are discussed in Paisley & Carin (2009); Chen et al. (2011; 2013). While the model may look somewhat complicated, local conjugacy admits Gibbs sampling or variational Bayes inference (Chen et al., 2011; 2013).

In Chen et al. (2011; 2013) a deep model was developed based on (1), by using $\mathbf{S}^{(n,k)} \stackrel{\text{def}}{=} \mathbf{Z}^{(n,k)} \odot \mathbf{W}^{(n,k)}$ as the input of the layer above. In order to do this, a pooling operation (*e.g.*, the max-pooling used in Chen et al. (2013)) is employed, reducing the feature dimension as one moves to higher layers. However, the model was learned by stacking layers upon each other, without subsequent overall refinement. This was because use of deterministic max pooling undermined development of a proper top-down generative model that coupled all layers; therefore, in Chen et al. (2013) the model in (1) was used sequentially from bottom-up, but the overall model parameters were never coupled when learning. To tackle this, we propose a *probabilistic pooling* procedure, yielding a top-down deep generative statistical structure, coupling all parameters when performing learning. As discussed when presenting results, this joint learning of all layers plays a critical role in improving model performance. The stochastic pooling applied here is closely related to that in Zeiler & Fergus (2013); Lee et al. (2009).

2.2 PRETRAINING & STOCHASTIC POOLING

Parameters of the deep model are learned by first analyzing one layer of the model at a time, starting at the bottom layer (touching the data), and sequentially stacking layers. The parameters of each layer of the model are learned separately, conditioned on parameters of the layers learned thus far (like in Chen et al. (2011; 2013)). The parameters learned in this manner serve as *initializations* for the top-down refinement step, discussed in Sec. 2.3, in which parameters at all layers of the deep model are learned jointly.

Assume an L-layer model, with layer L the top layer, and layer 1 at the bottom, closest to the data. In the pretraining stage, the output of layer l is the input to layer $l+1$, after pooling. Layer $l \in \{1, \ldots, L\}$ has K_l dictionary elements, and we have:

$$\mathbf{X}^{(n,l+1)} = \sum_{k_{l+1}=1}^{K_{l+1}} \mathbf{D}^{(k_{l+1},l+1)} * \left(\mathbf{Z}^{(n,k_{l+1},l+1)} \odot \mathbf{W}^{(n,k_{l+1},l+1)} \right) + \mathbf{E}^{(n,l+1)} \tag{5}$$

$$\mathbf{X}^{(n,l)} = \sum_{k_l=1}^{K_l} \mathbf{D}^{(k_l,l)} * \underbrace{\left(\mathbf{Z}^{(n,k_l,l)} \odot \mathbf{W}^{(n,k_l,l)} \right)}_{=\mathbf{S}^{(n,k_l,l)}} + \mathbf{E}^{(n,l)} \tag{6}$$

The expression $\mathbf{S}^{(n,k_l,l)}$ is a 2D (spatial) activation map, for image n, model layer l, dictionary element k_l. The expression $\mathbf{X}^{(n,l+1)}$ may be viewed as a 3D entity, with its k_l-th plane defined by a "pooled" version of $\mathbf{S}^{(n,k_l,l)}$ (pooling discussed next). The dictionary elements $\mathbf{D}^{(k_l,l)}$ and residual $\mathbf{E}^{(n,l)}$ are also three dimensional (each 2D plane of $\mathbf{D}^{(k_l,l)}$ and $\mathbf{E}^{(n,l)}$ is the spatial-dependent structure of the corresponding features), and the convolution is performed in the 2D spatial domain, simultaneously for each layer of the feature map.

We now discuss the relationship between $\mathbf{S}^{(n,k_l,l)}$ and layer k_l of $\mathbf{X}^{(n,l+1)}$. The 2D activation map $\mathbf{S}^{(n,k_l,l)}$ is partitioned into $n_x \times n_y$ dimensional contiguous blocks (pooling blocks with respect to layer $l+1$ of the model); see the left part of Figure 1. Associated with each block of pixels in $\mathbf{S}^{(n,k_l,l)}$ is one pixel at layer k_l of $\mathbf{X}^{(n,l+1)}$; the relative locations of the pixels in $\mathbf{X}^{(n,l+1)}$ are the same as the relative locations of the blocks in $\mathbf{S}^{(n,k_l,l)}$. Within each block of $\mathbf{S}^{(n,k_l,l)}$, either all $n_x n_y$ pixels are zero, or only one pixel is non-zero, with the position of that pixel selected stochastically via a multinomial distribution. Each pixel at layer k_l of $\mathbf{X}^{(n,l+1)}$ equals the largest-amplitude element in the associated block of $\mathbf{S}^{(n,k_l,l)}$ ($i.e.$, max pooling). Hence, if all elements of a block of $\mathbf{S}^{(n,k_l,l)}$ are zero, the corresponding pixel in $\mathbf{X}^{(n,l+1)}$ is also zero. If a block of $\mathbf{S}^{(n,k_l,l)}$ has a (single) non-zero element, that non-zero element is the corresponding pixel value at the k_l-th layer of $\mathbf{X}^{(n,l+1)}$.

The bottom-up generative process for each block of $\mathbf{S}^{(n,k_l,l)}$ proceeds as follows (left part of Figure 1). The model first imposes that a given block of $\mathbf{S}^{(n,k_l,l)}$ is either all zero or has one non-zero element, and this binary question is modeled as the beta-Bernoulli representation of (6). If a given block has a non-zero value, the position of that value in the associated $n_x \times n_y$ block is defined by a multinomial distribution, and its value is modeled as $w_{i,j}^{(n,k_l,l)}$ represented in (6). The beta-Bernoulli step, followed by multinomial, are combined into one equivalent statistical representation, as discussed next.

Let $\mathbf{z}_{i',j'}^{(n,k_l,l)} \in \{0,1\}^{n_x n_y}$ denote the (i',j')-th block of $\mathbf{Z}^{(n,k_l,l)}$ at layer l, where $i' = 1, \ldots, \frac{N_x}{n_x}; j' = 1, \ldots, \frac{N_y}{n_y}$ assuming integer divisions. We introduce a latent variable $\mathbf{c}_{i',j'}^{(n,k_l,l)} \in \{0,1\}^{n_x n_y + 1}$ to implement at most one non-zero

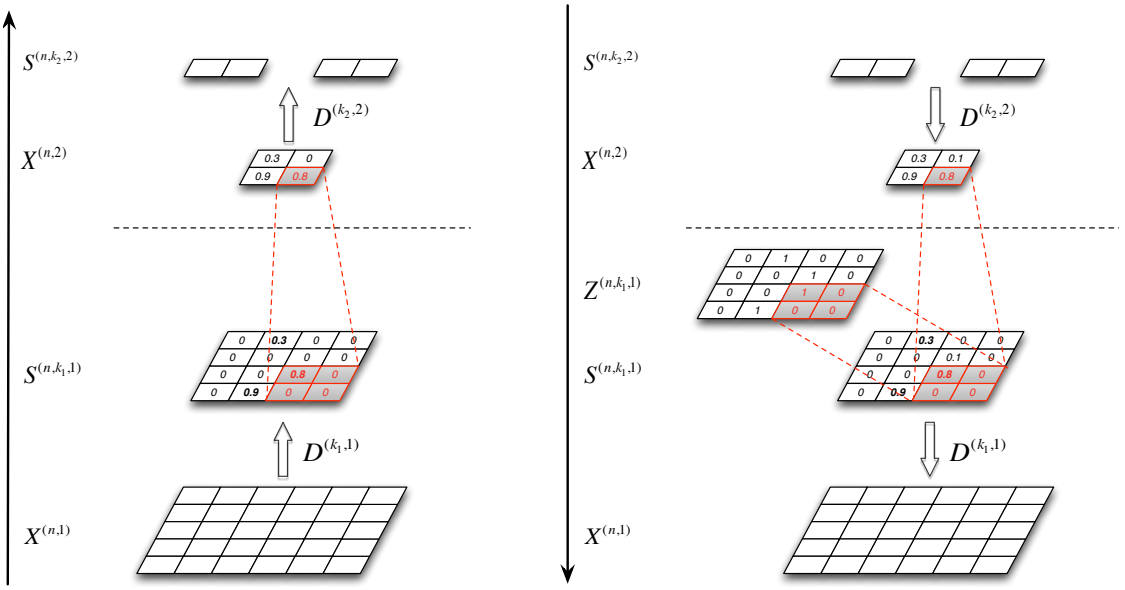

Figure 1: Schematic of the proposed generative process. Left: bottom-up pretraining, right: top-down refinement. (Zoom-in for best visulization and a larger version can be found in the Supplementary Material.)

element out of the $n_x n_y$ entries in $\{z^{(n,k_l,l)}_{i',j',m}\}_{m=1}^{n_x n_y}$ through

$$z^{(n,k_l,l)}_{i',j',m} = c^{(n,k_l,l)}_{i',j',m}, \qquad \mathbf{c}^{(n,k_l,l)}_{i',j'} \sim \text{Mult}(1; \boldsymbol{\theta}^{(n,k_l,l)}), \qquad \boldsymbol{\theta}^{(n,k_l,l)} \sim \text{Dir}\left(\frac{1}{n_x n_y + 1}\right), \tag{7}$$

where $\text{Mult}(\cdot)$ and $\text{Dir}(\cdot)$ denote multinomial and Dirichlet distribution, respectively (the Dirichlet distribution has a *set* of parameters, and here we imply that are equal, and set to the value indicated in $\text{Dir}(\cdot)$). $\mathbf{c}^{(n,k_l,l)}_{i',j'}$ has $(n_x n_y + 1)$ entries, of which only one is equal to 1. If the last element is 1, this means all $\{z^{(n,k_l,l)}_{i',j',m}\}_{m=1}^{n_x n_y} = 0$. Since the (i',j')-th block at layer l corresponds to one element at layer $(l+1)$, we have

$$s^{(n,k_l,l)}_{i',j',m} = x^{(n,k_l,l+1)}_{i',j'} z^{(n,k_l,l)}_{i',j',m}, \quad \forall m = 1, \ldots, n_x n_y \tag{8}$$

Hence, if the last element of $\mathbf{c}^{(n,k_l,l)}_{i',j'}$ is 1, all elements of block (i',j') are zero; if not, the location of the non-zero element in the first $n_x n_y$ entries of $\mathbf{c}^{(n,k_l,l)}_{i',j'}$ locates the position of the non-zero element in the corresponding block. The remaining parts of the model are represented as in (6).

In the pretraining phase, we start with $\mathbf{X}^{(n,1)}$, which is the data $\mathbf{X}^{(n)}$. We learn $\{\mathbf{S}^{(n,k_1,1)}\}_{k_1=1,K_1}$ using the blocked activation weights, via Gibbs sampling, where the multinomial distribution associates each non-zero element with a position in the corresponding block. The MAP Gibbs sample is then selected, defining model parameters for the layer under analysis. The "stacked" and pooled $\{\mathbf{S}^{(n,k_1,1)}\}_{k_1=1,K_1}$ are used to define $\mathbf{X}^{(n,2)}$, and the learning procedure then continues, learning dictionary elements $\mathbf{D}^{(k_2,2)}$ and activation maps $\{\mathbf{S}^{(n,k_2,2)}\}_{k_2=1,K_2}$, again via Gibbs sampling and MAP selection. This continues sequentially up to the L-th, or top, layer. For the top layer, since no pooling is necessary, the beta-Bernoulli prior in (2) is used.

2.3 Model Refinement With Stochastic Pooling

The learning performed with the top-down generative model (right part of Fig. 1) constitutes a *refinement* of the parameters learned during pretraining, and the excellent initialization constituted by the parameters learned during pretraining is key to the subsequent model performance.

In the refinement phase, the equations are (almost) the same, but we now proceed top down, from (5) to (6). The generative process constitutes $\mathbf{D}^{(k_{l+1},l+1)}$ and $\mathbf{Z}^{(n,k_{l+1},l+1)} \odot \mathbf{W}^{(n,k_{l+1},l+1)}$, and after convolution $\mathbf{X}^{(n,l+1)}$ is manifested; the $\mathbf{E}^{(n,l)}$ is now absent at all layers, except layer $l = 1$, at which the fit to the data is performed. Each element of $\mathbf{X}^{(n,l+1)}$ has an associated pooling *block* in $\mathbf{S}^{(n,k_l,l)}$. Via a multinomial distribution like in pretraining, each element of $\mathbf{X}^{(n,l+1)}$ is mapped to one position in the corresponding block of $\mathbf{S}^{(n,k_l,l)}$, and all other elements

in that $n_x \times n_y$ block are set to zero. Since $\mathbf{X}^{(n,l+1)}$ is manifested top-down as a convolution of $\mathbf{D}^{(k_{l+1},l+1)}$ and $\mathbf{Z}^{(n,k_{l+1},l+1)} \odot \mathbf{W}^{(n,k_{l+1},l+1)}$, $\mathbf{X}^{(n,l+1)}$ will in general have no elements exactly equal to zero (but many will be small, based on the pretraining). Hence, *each* block of $\mathbf{S}^{(n,k_l,l)}$ will have one non-zero element, with position defined by the multinomial[1].

During pretraining many blocks of $\mathbf{S}^{(n,k_l,l)}$ will be all-zero since we preferred a sparse representation, while during refinement this sparsity requirement is relaxed, and in general each pooling block of $\mathbf{S}^{(n,k_l,l)}$ will have one non-zero element (but it is still sparse), and this value is mapped via pooling to the corresponding pixel in $\mathbf{X}^{(n,l)}$. In pretraining the Dirichlet and multinomial distributions were of size $n_x n_y + 1$, allowing the all-zero activation block; during refinement the multinomial and Dirichlet are of dimensions $n_x n_y$. The corresponding $n_x n_y$ Dirichlet and multinomial parameters from pretraining are used to constitute initializations for refinement.

2.4 Top-Level Features and Testing

In order to understand deep convolutional models, researchers have visualized dictionary elements mapped to the image level (Zeiler & Fergus, 2014). One key challenge of this visualization is that one dictionary element at high layers can have *multiple* representations at the layer below, given different activations in each pooling block (in our model, this is manifested by the stochasticity associated with the multinomial-based pooling). Zeiler & Fergus (2014) showed different versions of the same upper-layer dictionary element at the image level. Because of this capability of accurate dictionary localization at each layer, deep convolutional models perform well in classification. However, also due to these multiple representations, during testing, one has to infer dictionary activations layer by layer (via deconvolution), which is computationally expensive. In order to alleviate this issue, Kavukcuoglu et al. (2010) proposed an approximation method using convolutional filter banks (fast because there is no explicit deconvolution) followed by a nonlinear function. Though efficient at test time, in the training step one must simultaneously learn deconvolutional dictionaries and associated filterbanks, and the choice of non-linear function is critical to the performance of the model. Moreover, in the context of the framework proposed here, it is difficult to integrate the approach of Kavukcuoglu et al. (2008; 2010) into a Bayesian model.

We propose a new approach to accelerate testing. After performing model learning (after refinement), we project top-layer dictionary elements down to the data plane. At test, deconvolution is only performed once, using the top-layer dictionary elements mapped to the data plane. The top-layer activation strengths inferred via this deconvolution are then used in a subsequent classifier. The different manifestations of a top-layer dictionary element mapped to the data plane are constituted by different (stochastic) pooling mappings via the multinomial. To select top-layer dictionary elements in the data plane, used for test, we employ maximum-likelihood (ML) dictionary elements, with ML performed across the different choices of the max pooling at each layer. Hence, after this ML-based top-layer dictionary selection, a pixel at layer $l+1$ is mapped to the same location in the associated layer l block, for all convolutional shifts (same max-pooling map for all shifts at a given layer). Hence, the key approximation is that the stochastic pooling employed for each pixel at layer $l+1$ to a position in a block at layer l is replaced by an ML-based *deterministic* pooling (possibly a different deterministic map at each layer). This simple approach has the advantage of Zeiler & Fergus (2014) at test, in that we retain the deconvolution operation (unlike Kavukcuoglu et al. (2010)), but deconvolution must only be performed *once* (not at each layer). In the experiments presented below, when visualizing inferred dictionary elements in the image plane, this ML-based dictionary selection is employed. More details on this aspect of the model are provided in the Supplementary Material.

3 Gibbs-Sampling-Based Learning and Inference

Due to local conjugacy at every component of the model, the local conditional posterior distribution for all parameters of our model is manifested in closed form, yielding efficient Gibbs sampling (see Supplementary Material for details). As in all previous convolutional models of this type, the FFT is leveraged to accelerate computation of the convolution operations, here within Gibbs update equations.

In the pre-training step, we select the ML sample from 500 collection samples, after first computing and discarding 1500 burn-in samples. The same number of burn-in and collection samples, with ML selection, is performed for model refinement. This ML selection of collection samples shares the same spirit as Geman & Geman (1984), in the sense of yielding a MAP solution (*not* attempting to approximate the full posterior). During testing, we select the ML sample across 200 deconvolutional samples, after first discarding 500 burn-in samples.

[1]We also considered a model exactly like in pretraining, which in the pooling step a pixel in $\mathbf{X}^{(n,l+1)}$ could be mapped via the multinomial to an all-zero activation block in layer l; the results are essentially unchanged from the method discussed above.

4 EXPERIMENTAL RESULTS

We here apply our model to the MNIST and Caltech 101 datasets. We compare dictionaries (viewed in the data plane) before and after refinement. Classification results (average of 10 trials) using top-layer features are presented for both datasets. As in (Paisley & Carin, 2009), the hyperparameters are set as $a_0 = 1/K, b_0 = 1 - 1/K$, where K is the number of dictionary elements at the corresponding layer, and $a_w = b_w = a_d = b_d = a_e = b_e = 10^{-6}$; these are standard hyperparameter settings (Paisley & Carin, 2009) for such models, and no tuning or optimization was performed. All code is written in MATLAB and executed on a desktop with 3.8 GHz CPU and 24G memory. Model training including refinement with one class (30 images) of Caltech 101 takes about 40 CPU minutes, and testing (deconvolution) for one image takes less than 1 second. These results were run on a single computer, for demonstration, and acceleration via parallel implementation, GPUs (Krizhevsky et al., 2012), and coding in C will be considered in the future; the successes realized recently in accelerating convolution-based models of this type are transferrable to our model.

MNIST Dataset We first consider the widely studied MNIST data (`http://yann.lecun.com/exdb/mnist/`), which has 60,000 training and 10,000 testing images, each 28×28, for digits 0 through 9. A two layer model is used with dictionary size 8×8 ($n_{d_x} = n_{d_y} = 8$) at the first layer and 6×6 at the second layer; the pooling size is 3×3 ($n_x = n_y = 3$) and the number of dictionary elements at layers 1 and 2 are $K_1 = 32$ and $K_2 = 160$, respectively. We obtained these number of dictionary elements via setting the initial dictionary number to a relatively large value in the pre-training step and discarding infrequently used elements by counting the corresponding binary indicator $\mathbf{Z}$, *i.e.*, *inferring* the number of needed dictionary elements, as in Chen et al. (2013).

Table 1: Classification Error of MNIST data

Methods	Test error
DBN Hinton & Salakhutdinov (2006)	1.20%
CBDN Lee et al. (2009)	0.82%
2-layer Conv. Net + 2-layer Classifier Jarrett et al. (2009)	0.53%
6-layer Conv. Net + 2-layer Classifier + elastic distortions Ciresan et al. (2011)	0.35%
MCDNN Ciresan et al. (2012)	0.23%
SPCNN Zeiler & Fergus (2013)	
Average Pooling	0.83%
Max Pooling	0.55%
Stochastic Pooling	0.47%
HBP Chen et al. (2013), 2-layer cFA + 2-layer features	
MCMC (10000 Training)	0.89%
Batch VB (10000 Training)	0.95%
online VB (60000 Training)	0.96%
Ours, 2-layer model + 1-layer features	
60000 Training	0.42%
10000 Training	0.68%
5000 Training	1.02%
2000 Training	1.11%
1000 Training	1.66%

Table 1 summaries the classification results of our model compared with some related results, on the MNIST data. The second (top) layer features corresponding to the refined dictionary are sent to a nonlinear support vector machine (SVM) (Chang & Lin, 2011) with Gaussian kernel, in a one-vs-all multi-class classifier, with classifier parameters tuned via 5-fold cross-validation (no tuning on the deep feature learning). Rather than concatenating features at all layers as in Zeiler & Fergus (2013); Chen et al. (2013), we only use the top layer features as the input to the SVM (deconvolution is only performed with top-layer dictionary elements), which saves much computation time (as well as memory) in both inference and classification, since the feature size is small. When the model is trained using all 60000 digits, we achieve an error rate of 0.42% on testing, which is very close to the state-of-the-art, but with a relatively simpler model compared to Ciresan et al. (2012); the error rate obtained using features learned after pretraining, before refinement, are similar to those in Chen et al. (2013) (0.9% error), underscoring the importance of the refinement step. We further plot the testing error in Fig. 2c (bottom part) when the training size is reduced compared to the results reported in Zeiler & Fergus (2013). It can be seen that our model outperforms every approach in Zeiler & Fergus (2013). In order to examine the properties of the learned model, in Fig. 2a we visualize trained dictionaries at layer 2 mapped down to the data level. It is observed qualitatively that refinement improves the dictionary; the atoms after refinement are much sharper. If the average pooling described in Zeiler & Fergus (2013) is used, the dictionaries are blurry (middle-left part of Fig. 2a). When a threshold is imposed on the refined dictionary elements, they look like digits (rightmost part).

To further verify the efficacy of our model, we show in Fig. 2b the interpolation results of digits with half missing, as in Lee et al. (2009). A one-layer model cannot recover the digits, while a two-layer model provides a good recovery (bottom row of Fig. 2b). Furthermore, by using our refinement approach, the recovery is much clearer (comparing the bottom-left part and bottom-middle part of Fig. 2b). Given this excellent performance, more challenging interpolation results are shown in Fig. 2c (upper part), where we cannot identify any digits from the observations; even in this case, the model can provide promising reconstructions.

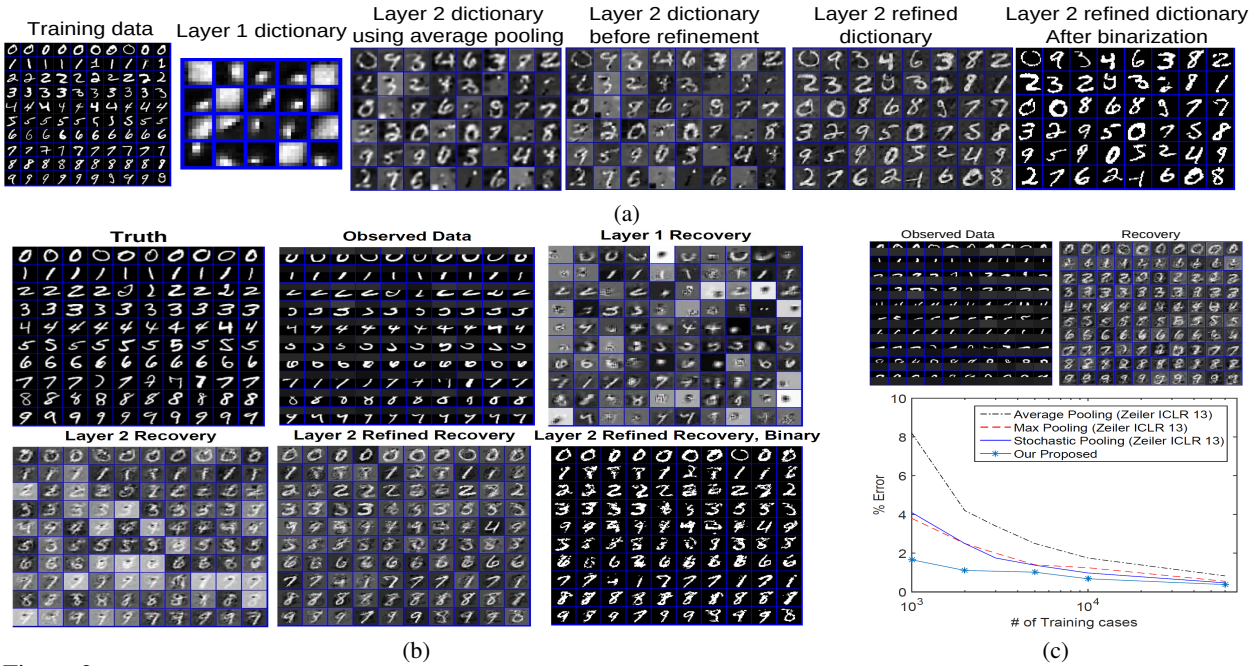

Figure 2: (a) Visualization of the dictionary learned by the proposed model. Note the refined dictionary (right) is much sharper than the dictionary before refinement (middle). (b) Missing data interpolation of digits. (c) Upper part: a more challenging case for missing data interpolation of digits. Bottom part: testing error when training with reduced dataset sizes on MNIST.

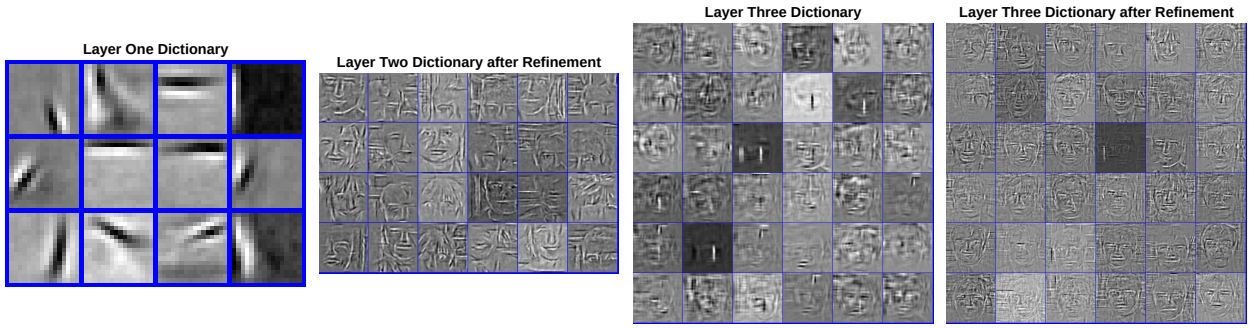
Layer One Dictionary
Layer Two Dictionary after Refinement
Layer Three Dictionary
Layer Three Dictionary after Refinement

Figure 3: Dictionary elements in each layer trained with 64 "face easy" images from Caltech 101.

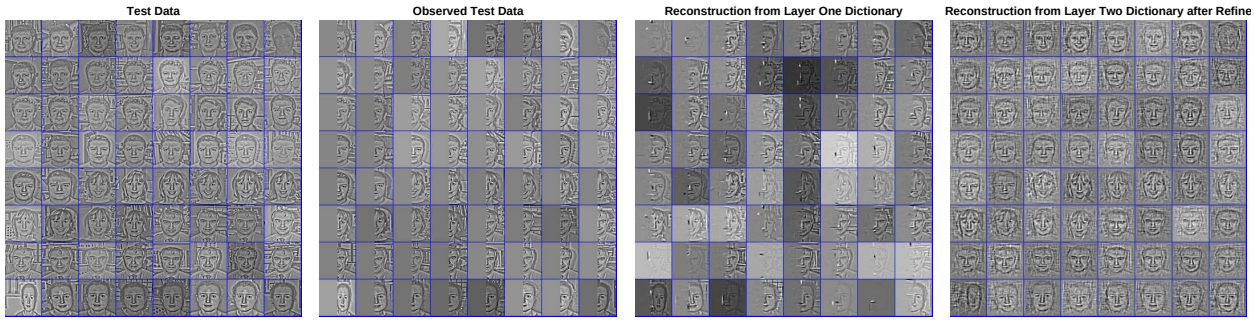
Test Data
Observed Test Data
Reconstruction from Layer One Dictionary
Reconstruction from Layer Two Dictionary after Refine

Figure 4: Face data interpolation using a 2-layer model. From left to right: truth, observed data, layer-1 recovery, layer-2 recovery.

Caltech 101 Dataset We next consider the Caltech 101 dataset. First we analyze our model with images in the "easy face" category; 64 images (after local contrast normalization (Jarrett et al., 2009)) have been resized to 128×128 and a three-layer deep model is used. At layers 1, 2 and 3, the number of dictionary elements is set respectively to $K_1 = 16$, $K_2 = 24$ and $K_3 = 36$ (these inferred in the pretraining step, as discussed above), with dictionary sizes 17×17, 9×9 and 6×6. The pooling sizes are 4×4 (layer 1 to layer 2) and 2×2 (layer 2 to layer 3). Example learned dictionary elements are mapped to the image level and shown in Fig. 3. It can be seen that the first-layer dictionary extracts edges of the images, while the second-layer dictionary elements look like a part of the face and the third-layer elements are almost entire faces. We can see the improvement manifested by refinement by comparing the right two parts in Fig. 3 (the dictionaries after refinement are sharper). Similar to the MNIST example, we also show in Fig. 4 the interpolation results of face data with half missing, using a two-layer model (the dictionary sizes are 14×14 and 13×13 at layers 1 and 2, respectively, with max-pooling size 3×3.). It can be seen the missing parts are recovered progressively more accurately considering a one- and two-layer model. Though the background is a little noisy, each face is recovered in great detail by the second layer dictionary (a three-layer model gives similar results, omitted here for brevity).

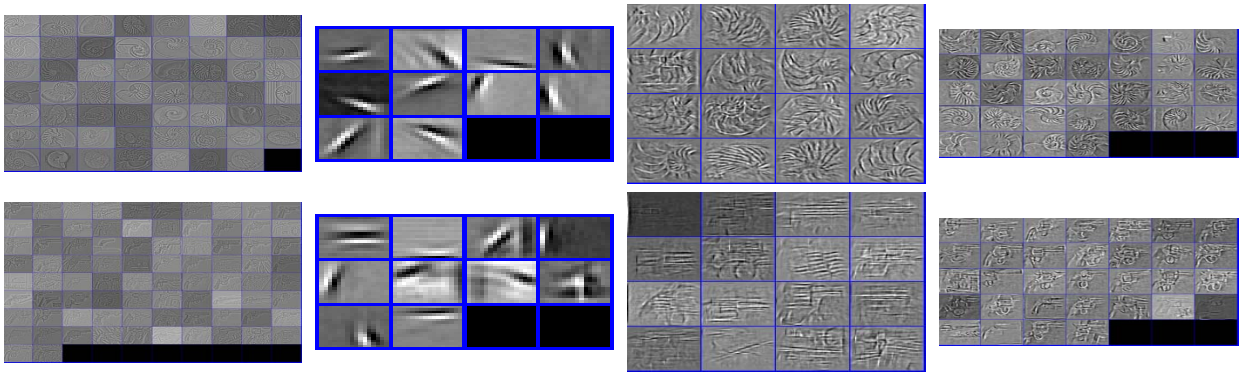

Figure 5: Trained dictionaries per class mapped to the data plane. Row 1-2: nautilus, revolver. Column 1-4: training images after local contrast normalization, layer-1 dictionary, layer-2 dictionary, layer-3 dictionary.

We develop Caltech 101 dictionaries by learning on each data class in isolation, and then concatenate all (top-layer) dictionaries when learning the classifier. In Figure 5 we depict dictionary elements learned for two data classes, projected to the image level (more results are shown in the Supplementary Material). It can be seen the layer-1 dictionary elements are similar for the two data classes, while the upper-layer dictionary elements are data-class dependent. One problem of this parallel training is that the dictionary may be redundant across image classes (especially at the first layer). However, during testing, using the proposed approach, we only use top-layer dictionaries, which are typically distinct across data classes (for the data considered).

For Caltech 101 classification, we follow the setup in Yang et al. (2009), selecting 15 and 30 images per category for training, and testing on the rest. The features of testing images are inferred based on the top-layer dictionaries and sent to a multi-class SVM; we again use a Gaussian kernel non-linear SVM with parameters tuned via cross-validation. Ours and related results are summarized in Table 2. For our model, we present results based on 2-layer and 3-layer models. It can be seen that our model (the 3-layer one) provides results close to the state-of-the-art in Zeiler & Fergus (2014), which used a much more complicated model (*i.e.*, a 7-layer convolutional network and used the ImageNet dataset to pretrain the network), and our results are also very close to the state-of-the-art results using hand-crafted features (*e.g.*, SIFT in Li et al. (2013)). Based on features learned by our model at the pretraining stage, our classification performance is comparable to that of the HBP model in Chen et al. (2013) (around 65% accuracy for a 2-layer model, when training with 30 examples per class), with our results demonstrating a 17% improvement in performance after model refinement.

Table 2: Classification Accuracy Rate of Caltech-101.

# Training Images per Category	15	30
DN Zeiler et al. (2010)	58.6 %	66.9%
CBDN Lee et al. (2009)	57.7 %	65.4%
HBP Chen et al. (2013)	58%	65.7%
ScSPM Yang et al. (2009)	67 %	73.2%
P-FV Seidenari et al. (2014)	71.47%	80.13%
R-KSVD Li et al. (2013)	79 %	83%
Convnet Zeiler & Fergus (2014)	83.8 %	86.5%
Ours, 2-layer model + 1-layer features	70.02%	80.31%
Ours, 3-layer model + 1-layer features	75.24%	82.78%

5 CONCLUSIONS

A deep generative convolutional dictionary-learning model has been developed within a Bayesian setting, with efficient Gibbs-sampling-based MAP parameter estimation. The proposed framework enjoys efficient bottom-up and top-down probabilistic inference. A probabilistic pooling module has been integrated into the model, a key component to developing a principled top-down generative model, with efficient learning and inference. Extensive experimental results demonstrate the efficacy of the model to learn multi-layered features from images. A novel method has been developed to project the high-layer dictionary elements to the image level, and efficient single-layer deconvolutional inference is accomplished during testing. On the MNIST and Caltech 101 datasets, our results are very near the state of the art, but with relatively simple model complexity at test. Future work includes performing deep feature learning and classifier design jointly. The algorithm will also be ported to a GPU-based implementation, allowing scaling to large-scale datasets.

REFERENCES

Chang, C.-C. and Lin, C.-J. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2011.

Chen, B., Polatkan, G., Sapiro, G., Carin, L., and Dunson, D. B. The hierarchical beta process for convolutional factor analysis and deep learning. In *ICML*, 2011.

Chen, B., Polatkan, G., Sapiro, G., Blei, D., Dunson, D., and Carin, L. Deep learning with hierarchical convolutional factor analysis. *IEEE T-PAMI*, 2013.

Ciresan, D., Meier, U., and Schmidhuber, J. Multi-column deep neural networks for image classification. In *CVPR*, 2012.

Ciresan, D. C., Meier, U., Masci, J., Gambardella, L. M., and Flexible, J. Schmidhuber. high performance convolutional neural networks for image classification. *IJCAI*, 2011.

Geman, S. and Geman, D. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE T-PAMI*, 1984.

Hinton, G. and Salakhutdinov, R. Reducing the dimensionality of data with neural networks. *Science*, 2006.

Jarrett, K., Kavukcuoglu, K., Ranzato, M., and LeCun, Y. What is the best multi-stage architecture for object recognition? *ICCV*, 2009.

Kavukcuoglu, K., Ranzato, M., and LeCun, Y. Fast inference in sparse coding algorithms with applications to object recognition. In *ArXiv 1010.3467*, 2008.

Kavukcuoglu, K., Sermanet, P., Boureau, Y-L., Gregor, K., Mathieu, M., and LeCun, Y. Learning convolutional feature hierarchies for visual recognition. *NIPS*, 2010.

Krizhevsky, A., S., Ilya, and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, 1998.

Lee, H., Grosse, R., Ranganath, R., and Ng, A. Y. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. *ICML*, 2009.

Li, Q., Zhang, H., Guo, J., Bhanu, B., and An, L. Reference-based scheme combined with K-svd for scene image categorization. *IEEE Signal Processing Letters*, 2013.

Paisley, J. and Carin, L. Nonparametric factor analysis with beta process priors. In *ICML*, 2009.

Seidenari, L., Serra, G., Bagdanov, A., and Del Bimbo, A. Local pyramidal descriptors for image recognition. *IEEE T-PAMI*, 2014.

Yang, J., Yu, K., Gong, Y., and Huang, T. Linear spatial pyramid matching using sparse coding for image classification. In *CVPR*, 2009.

Zeiler, M. and Fergus, R. Stochastic pooling for regularization of deep convolutional neural networks. *ICLR*, 2013.

Zeiler, M. and Fergus, R. Visualizing and understanding convolutional networks. *ECCV*, 2014.

Zeiler, M., Kirshnan, D., Taylor, G., and Fergus, R. Deconvolutional networks. *CVPR*, 2010.

Supplementary Material

6 Conditional Posteriori Distributions for Gibbs Sampling

In the $l-$th layer, the model can be formed as:

$$\mathbf{X}^{(n,k_{l-1},l)} = \sum_{k_l=1}^{K_l} \mathbf{D}^{(k_l,l)} * \left(\mathbf{Z}^{(n,k_l,l)} \odot \mathbf{W}^{(n,k_l,l)}\right) + \mathbf{E}^{(n,k_{l-1},l)} \tag{9}$$

For simplification, we define the following symbols (operations):

$$X^{-(n,k_{l-1},l)} = X^{(n,k_{l-1},l)} - \sum_{k_l} D^{(k_{l-1},k_l,l)} * (Z^{(n,k_l,l)} \odot W^{(n,k_l,l)}) \tag{10}$$

$$X_{-k_l}^{(n,k_{l-1},l)} = X^{(n,k_{l-1},l)} - \sum_{t\neq k_l} D^{(k_{l-1},t,l)} * (Z^{(n,t,l)} \odot W^{(n,t,l)}) \tag{11}$$

$$(D^{(k_{l-1},k_l,l)})^2 = D^{(k_{l-1},k_l,l)} \odot D^{(k_{l-1},k_l,l)} \tag{12}$$

$$(W^{(n,k_l,l)})^2 = W^{(n,k_l,l)} \odot W^{(n,k_l,l)} \tag{13}$$

$$(X^{-(n,k_{l-1},l)})^2 = X^{-(n,k_{l-1},l)} \odot X^{-(n,k_{l-1},l)} \tag{14}$$

The symbol $\odot$ is the element-wise product operator and $\oslash$ is the element-wise division operator.

$A = B \circledast C$ means if $B \in R^{n_B \times n_B}$ and $C \in R^{n_C \times n_C}$, $A \in R^{(n_B-n_C+1)\times(n_B-n_C+1)}$ with element (i,j)

$$A_{i,j} = \sum_{p=1}^{n_C}\sum_{q=1}^{n_C} B_{p+i-1,q+j-1} C_{p,q}. \tag{15}$$

For each MCMC iteration, the samples are drawn from:

- Sample $D^{(k_{l-1},k_l,l)}$:

$$D^{(k_{l-1},k_l,l)}|- \sim N(\mu^{(k_{l-1},k_l,l)}, \Sigma^{(k_{l-1},k_l,l)}) \tag{16}$$

$$\Sigma^{(k_{l-1},k_l,l)} = 1 \oslash \left(\sum_{n=1}^{N} \gamma_e^{(n,k_{l-1},l)} \|Z^{(n,k_l,l)} \odot W^{(n,k_l,l)}\|_2^2 + \gamma_d^{(k_{l-1},k_l,l)}\right) \tag{17}$$

$$\hat{\mu}^{(k_{l-1},k_l,l)} = \Sigma^{(k_{l-1},k_l,l)} \odot \left[\sum_{n=1}^{N} \gamma_e^{(n,k_{l-1},l)} X_{-k_l}^{(n,k_{l-1},l)} \circledast (Z^{(n,k_l,l)} \odot W^{(n,k_l,l)})\right] \tag{18}$$

- Sample $\gamma_d^{(k_{l-1},k_l,l)}$:

$$\gamma_d^{(k_{l-1},k_l,l)}|- \sim \text{Gamma}\left(a_d + \frac{1}{2}, b_d + \frac{1}{2}(D^{(k_{l-1},k_l,l)})^2\right) \tag{19}$$

- Sample $W^{(n,k_l,l)}$:

$$W^{(n,k_l,l)}|- \sim N(\xi^{(n,k_l,l)}, \Lambda^{(n,k_l,l)}) \tag{20}$$

$$\Lambda^{(n,k_l,l)} = 1 \oslash \left(\sum_{k_{l-1}=1}^{K_{l-1}} \gamma_e^{(n,k_{l-1},l)} \|D^{(k_{l-1},k_l,l)}\|_2^2 Z^{(n,k_l,l)} + \gamma_W^{(n,k_l,l)}\right) \tag{21}$$

$$\xi^{(n,k_l,l)} = \Lambda^{(n,k_l,l)} \odot Z^{(n,k_l,l)} \odot \left[\sum_{k_{l-1}=1}^{K_{l-1}} \gamma_e^{(n,k_{l-1},l)} X_{-k_l}^{(n,k_{l-1},l)} \circledast D^{k_{l-1},k_l,l}\right] \tag{22}$$

- Sample $\gamma_W^{(n,k_l,l)}$

$$\gamma_W^{(n,k_l,l)}|- \sim \text{Gamma}\left(a_W + \frac{1}{2}, b_W + \frac{1}{2}(W^{(n,k_l,l)})^2\right) \tag{23}$$

- Sample $Z_{i',j'}^{(n,k_l,l)}$:

  Let $m = 1, ..., n_x n_y$, $i = 1, ..., N_x$, $j = 1, ..., N_y$; $i' = 1, ..., N_x/n_x$, $j = 1, ..., N_y/n_x$; we can find (i', j', m) and (i, j) are one-to-one correspondence. From

$$Y^{(n,k_l,l)} = \sum_{k_{l-1}=1}^{K_{l-1}} \gamma_e^{(n,k_{l-1},l)} \left[\|D^{(k_{l-1},k_l,l)}\|_2^2 \odot \left(W^{(n,k_l,l)}\right)^2 - 2\left(X_{-k_l}^{(n,k_{l-1},l)} \circledast D^{k_{l-1},k_l,l}\right) \odot W^{(n,k_l,l)} \right], \quad (24)$$

  and

$$\hat{\theta}_{i',j',m}^{(n,k_l,l)} = \theta_m^{(n,k_l,l)} Y_{i',j',m}^{(n,k_l,l)} \qquad \text{for } m = 1, ..., n_x n_y, \quad (25)$$

  we have

$$P(Z_{i',j',m}^{(n,k_l,l)} = 1) = \frac{\hat{\theta}_{i',j',m}^{(n,k_l,l)}}{\sum_{t=1(\neq m)}^{n_x n_y} \hat{\theta}_{i',j',m}^{(n,k_l,l)} + \theta_{n_x n_y+1}^{(n,k_l,l)}}. \quad (26)$$

- Sample $\theta^{(n,k_l,l)}$:

$$\theta^{(n,k_l,l)}|- \sim \text{Dirichlet}(\alpha^{(n,k_l,l)}) \quad (27)$$

$$\alpha_m^{(n,k_l,l)} = \frac{1}{n_x n_y + 1} + \sum_{i'} \sum_{j'} z_{i',j',m}^{(n,k_l,l)} \qquad \text{for} \quad m = 1, ..., n_x n_y, \quad (28)$$

$$\alpha_{n_x n_y+1}^{(n,k_l,l)} = \frac{1}{n_x n_y + 1} + \sum_{i'} \sum_{j'} (1 - \sum_{m} z_{i',j',m}^{(n,k_l,l)}) \quad (29)$$

- Sample $\gamma_e^{(n,k_{l-1},l)}$:

$$\gamma_e^{(n,k_{l-1},l)}|- \sim \text{Gamma}\left(a_e + \frac{1}{2}, b_e + (X^{-(n,k_{l-1},l)})^2\right). \quad (30)$$

7 Projection of Dictionaries to the Data Layer

7.1 Notation

Assume $X \in \mathbb{R}^{N_x \times N_y}$ and $Y \in \mathbb{R}^{N_x/n_x \times N_y/n_y}$. Here $n_x, n_y \in N$ are the pooling ratio and the pooling map is $Z \in \{0,1\}^{N_x \times N_y}$. In the $(i', j')^{\text{th}}$ block of X and Z, there is at most one non-zero element, where $i' \in \{1, ..., \lceil N_x/n_x \rceil\}$, $j' \in \{1, ..., \lceil N_y/n_y \rceil\}$. Now, let $i \in \{1, ..., N_x\}$, $j \in \{1, ..., N_y\}$ then the following pooling and unpooling functions can be defined:

(1) Define $f : \mathbb{R}^{N_x \times N_y} \to \mathbb{R}^{N_x/n_x \times N_y/n_y}$, with $Y = f(X)$. Recall that within each pooling block, X has at most one non-zero element, and therefore

$$Y_{i',j'} = \sum_{s=1}^{n_x} \sum_{t=1}^{n_y} X_{(i'-1)n_x+s,(j'-1)n_y+t} \quad (31)$$

The following is an example to demonstrate f:

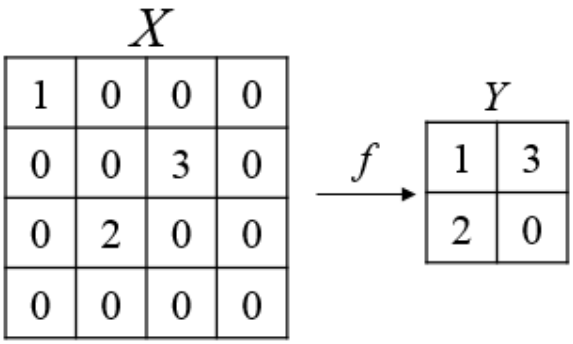

(2) Define $f^{-1}: \mathbb{R}^{N_x/n_x \times N_y/n_y} \times \{0,1\}^{N_x \times N_y} \to \mathbb{R}^{N_x \times N_y}$, with $X = f^{-1}(Y, Z)$

$$X_{i,j} = Y_{\lceil i/n_x \rceil, \lceil j/n_y \rceil} Z_{i,j} \tag{32}$$

The following is an example to demonstrate f^{-1}:

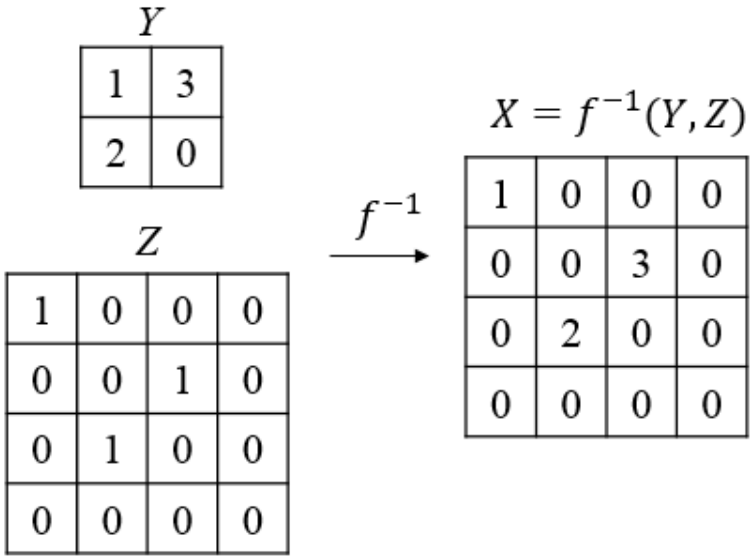

(3) Define $g: \mathbb{R}^{N_x/n_x \times N_y/n_y} \times \mathbb{N}^2 \to \mathbb{R}^{N_x \times N_y}$, with $A = g(Y, n_x, n_y)$

$$A_{i,j} = Y_{\left\lceil \frac{i}{n_x} \right\rceil, \left\lceil \frac{j}{n_y} \right\rceil}. \tag{33}$$

The following is an example to demonstrate g,with $n_x = n_y = 2$:

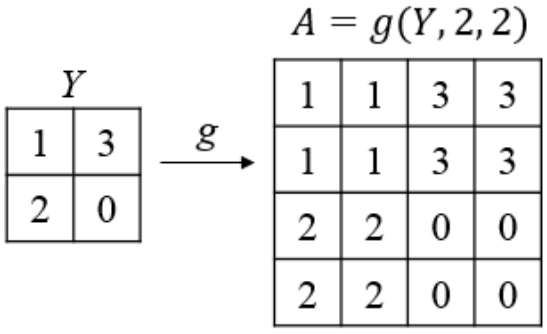

(4) Define $h: \mathbb{R}^{N_x/n_x \times N_y/n_y} \times \mathbb{N}^2 \to \mathbb{R}^{((N_x-1)n_x+1) \times ((N_x-1)n_y+1)}$, with $B = h(Y, n_x, n_y)$:

$$B_{i,j} = \begin{cases} Y_{\frac{i-1}{n_x}+1, \frac{j-1}{n_y}+1} & \text{if } \frac{i-1}{n_x} \text{ and } \frac{j-1}{n_y} \text{ both are integers} \\ 0 & \text{otherwise} \end{cases} \tag{34}$$

The following is an example to demonstrate h with $n_x = n_y = 2$:

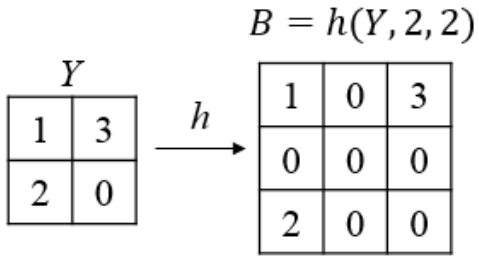

7.2 SOME USEFUL LEMMAS

Lemma 1. $f^{-1}(Y, Z) = g(Y) \odot Z$

Lemma 2. $g(g(Y, n_x^1, n_y^1), n_x^2, n_y^2) = g(Y, n_x^1 n_x^2, n_y^1 n_y^2)$

Lemma 3. $h(h(Y, n_x^1, n_y^1), n_x^2, n_y^2) = h(Y, n_x^1 n_x^2, n_y^1 n_y^2)$

Lemma 4. $g(D * W, n_x, n_y) = g(D, n_x, n_y) * h(W, n_x, n_y)$

Lemma 5. $h(D * W, n_x, n_y) = h(D, n_x, n_y) * h(W, n_x, n_y)$

The first three lemmas are obvious. Now, we provide the proof of lemma 4 and lemma 5 .
Lemma 4 proof:
Recall that the convolution operator $C = D * W$ means if $D \in \mathbb{R}^{n_{dx} \times n_{dy}}$ and $W \in \mathbb{R}^{n_{wx} \times n_{wy}}$, $C \in \mathbb{R}^{(n_{dx}+n_{wx}+1)\times(n_{dy}+n_{wy}+1)}$ then the $(i,j)^{\text{th}}$ element is given by

$$C_{i,j} = \sum_{p=1}^{n_{wx}} \sum_{q=1}^{n_{wy}} D_{i+n_{wx}-p, j+n_{wy}-q} W_{p,q} \tag{35}$$

where

$$D_{i,j} = 0 \qquad \text{if } i, j < 0 \qquad \text{or } i > n_{dx} \qquad \text{or } j > n_{dy}. \tag{36}$$

Let $\hat{C} = g(D * W, n_x, n_y)$, $\hat{D} = g(D, n_x, n_y)$, $\hat{W} = h(W, n_x, n_y)$ and $n'_{wx} = (n_{wx} - 1)n_x + 1$, $n'_{wy} = (n_{wy} - 1)n_y + 1$. We want to prove that $\hat{C} = \hat{D} * \hat{W}$. Deriving elementwise we have

$$\begin{aligned}
(\hat{D} * \hat{W})_{i',j'} &= \sum_{p'=1}^{n'_{wx}} \sum_{q'=1}^{n'_{wy}} \hat{D}_{i'+n_w-p', j'+n_w-q'} \hat{W}_{p',q'} \\
&= \sum_{p'=1}^{n'_{wx}} \sum_{q'=1}^{n'_{wy}} D_{\left\lceil \frac{i'+n'_{wx}-p'}{n_x} \right\rceil, \left\lceil \frac{j'+n'_{wy}-q'}{n_y} \right\rceil} \hat{W}_{p',q'} \\
&= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} D_{\left\lceil \frac{i'+(n_{wx}-1)n_x-(p-1)n_x}{n_x} \right\rceil, \left\lceil \frac{j'+(n_{wy}-1)n_y-(q-1)n_y}{n_y} \right\rceil} \hat{W}_{(p-1)n_x+1,(q-1)n_x+1} \\
&= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} D_{\left\lceil \frac{i'}{n_x} \right\rceil + n_{wx} - p, \left\lceil \frac{j'}{n_y} \right\rceil + n_{wy} - q} W_{p,q} \\
&= C_{\left\lceil \frac{i'}{n_x} \right\rceil, \left\lceil \frac{j'}{n_y} \right\rceil} \\
&= \hat{C}_{i',j'}.
\end{aligned} \tag{37}$$

Since if $p' = (p-1)n_x + 1$ and $q' = (q-1)n_x + 1$, then $\hat{W}_{p',q'} = W_{p,q}$, and if not $\hat{W}_{p',q'} = 0$.

Lemma 5 proof:
Let $\tilde{C} = h(D * W, n_x, n_y)$, $\tilde{D} = h(D, n_x, n_y)$, $\tilde{W} = h(W, n_x, n_y)$ and $n'_{wx} = (n_{wx} - 1)n_x + 1$, $n'_{wy} = (n_{wy} - 1)n_y + 1$.

We want to prove that $\tilde{C} = \tilde{D} * \tilde{W}$. Deriving elementwise we have:

$$\begin{aligned}(\tilde{D} * \tilde{W})_{i',j'} &= \sum_{p'=1}^{n'_{wx}} \sum_{q'=1}^{n'_{wy}} \tilde{D}_{i'+n_w-p',j'+n_w-q'} \tilde{W}_{p',q'} \\ &= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} \tilde{D}_{i'+(n_{wx}-p)n_x,j'+(n_{wy}-q)n_y} \tilde{W}_{(p-1)n_x+1,(q-1)n_x+1} \\ &= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} \tilde{D}_{i'+(n_{wx}-p)n_x,j'+(n_{wy}-q)n_y} W_{p,q} \end{aligned} \tag{38}$$

Therefore, when $\frac{i'-1}{n_x}$ and $\frac{i'-1}{n_y}$ both are integers

$$\begin{aligned}(\tilde{D} * \tilde{W})_{i',j'} &= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} \tilde{D}_{i'+(n_{wx}-p)n_x,j'+(n_{wy}-q)n_y} W_{p,q} \\ &= \sum_{p=1}^{n_w} \sum_{q=1}^{n_w} D_{\frac{i'-1}{n_x}+1+(n_{wx}-p),\frac{j'-1}{n_y}+1+(n_{wy}-q)} W_{p,q} \\ &= (D * W)_{\frac{i'-1}{n_x}+1,\frac{j'-1}{n_y}+1} \\ &= C_{\frac{i'-1}{n_x}+1,\frac{j'-1}{n_y}+1} \\ &= \tilde{C}_{i',j'} \end{aligned} \tag{39}$$

Otherwise, $(\tilde{D} * \tilde{W})_{i',j'} = 0 = \tilde{C}_{i',j'}$. Therefore, $\tilde{C} = \tilde{D} * \tilde{W}$

7.3 Layer Collapsing Approximation Method

Assume the spatial size of the l-th layer dictionary is $N^l_{dx} \times N^l_{dy}$ and the pooling ratio from the l-th layer to the $(l+1)$-th layer is $n^l_x \times n^l_y$. Then, the generative model can be formed as:

$$X^{(n,1)} = \sum_{k_1=1}^{K_1} D^{(k_1,1)} * f^{-1}(X^{(n,k_1,2)}, Z^{(n,k_1,1)}) + E^{(n)} \tag{40}$$

$$X^{(n,k_{l-1},l)} = \sum_{k_l=1}^{K_l} D^{(k_{l-1},k_l,l)} * f^{-1}(X^{(n,k_l,l+1)}, Z^{(n,k_l,l)}) \qquad \text{for} \quad l = 2, ..., L-1 \tag{41}$$

$$X^{(n,k_{L-1},L)} = \sum_{k_L=1}^{K_L} D^{(k_{L-1},k_L,L)} * (W^{(n,k_L,L)} \odot Z^{(n,k_L,L)}) \tag{42}$$

The dictionary can be projected down one layer using the following approximation

$$\begin{aligned} S^{(n,k_l,l)} &= f^{-1}(X^{(n,k_l,l+1)}, Z^{(n,k_l,l)}) \\ &= g(X^{(n,k_l,l+1)}, n^l_x, n^l_y) \odot Z^{(n,k_l,l)} \\ &= g\left(\sum_{k_{l+1}=1}^{K_{l+1}} D^{(k_l,k_{l+1},l+1)} * S^{(n,k_{l+1},l+1)}, n^l_x, n^l_y\right) \odot Z^{(n,k_l,l)} \\ &= \left[\sum_{k_{l+1}=1}^{K_{l+1}} g(D^{(k_l,k_{l+1},l+1)}, n^l_x, n^l_y) * h(S^{(n,k_{l+1},l+1)}, n^l_x, n^l_y)\right] \odot Z^{(n,k_l,l)} \\ &\approx \sum_{k_{l+1}=1}^{K_{l+1}} \left[g(D^{(k_l,k_{l+1},l+1)}, n^l_x, n^l_y) \odot \hat{Z}^{(n,k_l,l)}\right] * h(S^{(n,k_{l+1},l+1)}, n^l_x, n^l_y) \end{aligned} \tag{43}$$

where $\hat{Z}^{(n,k_l,l)} \in \{0,1\}^{N_{dx}^l n_x^l \times N_{dy}^l n_y^l}$. In each $n_x^l \times n_y^l$ block of $\hat{Z}^{(n,k_l,l)}$, there is at most one non-zero element. And down two layers by this approximation

$$
\begin{aligned}
S^{(n,k_{l-1},l-1)} &= g(\sum_{k_l=1}^{K_l} D^{(k_{l-1},k_l,l)} * S^{(n,k_l,l)}, n_x^{l-1}, n_y^{l-1}) \odot Z^{(n,k_{l-1},l-1)} \\
&\approx \sum_{k_l=1}^{K_l} g(D^{(k_{l-1},k_l,l)}) n_x^{l-1}, n_y^{l-1}) * h(S^{(n,k_l,l)}, n_x^{l-1}, n_y^{l-1}) \\
&= \sum_{k_l=1}^{K_l} g\left(D^{(k_{l-1},k_l,l)}, n_x^{l-1}, n_y^{l-1}\right) * h\left(\sum_{k_{l+1}=1}^{K_{l+1}} \left[g(D^{(k_l,k_{l+1},l+1)}, n_x^l, n_y^l) \odot \hat{Z}^{(n,k_l,l)}\right] * h(S^{(n,k_{l+1},l+1)}, n_x^l, n_y^l), n_x^{l-1}, n_y^{l-1}\right) \\
&= \sum_{k_l=1}^{K_l} g\left(D^{(k_{l-1},k_l,l)}, n_x^{l-1}, n_y^{l-1}\right) * \sum_{k_{l+1}=1}^{K_{l+1}} \left\{h\left[g(D^{(k_l,k_{l+1},l+1)}, n_x^l, n_y^l) \odot \hat{Z}^{(n,k_l,l)}, n_x^{l-1}, n_y^{l-1}\right] * h\left(S^{(n,k_{l+1},l+1)}, n_x^{l-1} n_x^l, n_y^{l-1} n_y^l\right)\right\} \\
&= \sum_{k_l=1}^{K_l} \sum_{k_{l+1}=1}^{K_{l+1}} g\left(D^{(k_{l-1},k_l,l)}, n_x^{l-1}, n_y^{l-1}\right) * h\left[g(D^{(k_l,k_{l+1},l+1)}, n_x^l, n_y^l) \odot \hat{Z}^{(n,k_l,l)}, n_x^{l-1}, n_y^{l-1}\right] * h\left(S^{(n,k_{l+1},l+1)}, n_x^{l-1} n_x^l, n_y^{l-1} n_y^l\right)
\end{aligned}
\tag{44}
$$

Iteratively collapsing all layers we arrive at the reduced model

$$
\begin{aligned}
X^{(n)} &= \sum_{k_1=1}^{K_1} D^{(k_1,1)} * \left[\sum_{k_2=1}^{K_2} \hat{D}^{(k_1,k_2,2)} * \left(\cdots * (\sum_{k_L=1}^{K_L} \hat{D}^{(k_{L-1},k_L,L)} * \hat{S}^{(n,k_L,L)})\right)\right] + E^{(n)} \\
&= \sum_{k_1=1}^{K_1} \sum_{k_2=1}^{K_2} \cdots \sum_{k_L=1}^{K_L} \left(D^{(k_1,1)} * \hat{D}^{(k_1,k_2,2)} * \cdots * \hat{D}^{(k_{L-1},k_L,L)} * \hat{S}^{(n,k_L,L)}\right) + E^{(n)} \\
&= \sum_{k_L=1}^{K_L} \left(\sum_{k_1=1}^{K_1} \cdots \sum_{k_{L-1}=1}^{K_{L-1}} D^{(k_1,1)} * \hat{D}^{(k_1,k_2,l)} * \cdots * \hat{D}^{(k_{L-1},k_L,L)}\right) * \hat{S}^{(n,k_L,L)} + E^{(n)} \\
&= \sum_{k_L=1}^{K_L} \tilde{D}^{(k_L)} * \hat{S}^{(n,k_L,L)} + E^{(n)}
\end{aligned}
\tag{45}
$$

where

$$
\hat{D}^{(k_{l-1},k_l,l)} = h\left(g(D^{(k_{l-1},k_l,l)}), \prod_{i=1}^{l-2} n_x^i, \prod_{i=1}^{l-2} n_y^i\right) \odot h\left(\tilde{Z}^{(n,k_{l-1},l)}, \prod_{i=1}^{l-2} n_x^i, \prod_{i=1}^{l-2} n_y^i\right), \qquad \text{for} \quad l = 3, ..., L-1 \tag{46}
$$

$$
\hat{D}^{(k_1,k_2,2)} = g(D^{k_1,k_2,2}, n_x^1, n_y^1) \odot \tilde{Z}^{(n,k_1,1)} \tag{47}
$$

$$
\hat{S}^{(n,k_L,L)} = h\left(S^{(n,k_L,L)}, \prod_{l=1}^{L-1} n_x^l, \prod_{l=1}^{L-1} n_y^l\right). \tag{48}
$$

and the collapsed (one-layer) dictionary is given by $\{\tilde{D}^{(k_L)}\}_{k_L=1}^{K_L}$.

8 TESTING

When given M testing images $\{X^{(m)}\}_{m=1}^{M}$, we learn the test feature from:

$$
X^{(m)} = \sum_{k_L=1}^{K_L} \tilde{D}^{k_L} * (W^{(m,k_L)} \odot Z^{(m,k_L)}) + E^{(m)} \tag{49}
$$

Let s_m be a vector "unfolded" from $\{S^{(m,k_L)} = W^{(m,k_L)} \odot Z^{m,k_L}\}_{k_L=1}^{K_L}$. s_m will be sent to a nonlinear support vector machine (SVM) with Gaussian kernel, in a one-vs-all multi-class classifier, with classifier parameters tuned via 5-fold cross-validation (no tuning on the deep feature learning).

Algorithm 1 Learning and Classification of the Deep Model

Require: Input images $\mathbf{X}$, model layers L, dictionary sizes $\{K_l\}_{l=1}^{L}$, pooling sizes n_x, n_y.

1: Model Training
   - Bottom-up pretraining with inference equations in Section 6.
   - Top-down refinement with equations in Section 6
2: Project top-layer dictionary into data level with equations in Section 7.
3: Testing with model in Section 8.
4: Features learned by (49) sent to SVM for classification.

9 Algorithm Flow-Chart

10 Figure 1

Figure 1 plots the large version of Figure 1 in the main paper.

11 Generated Images

We use the trained dictionary by MNIST and randomly generated weights to generate images shown in Figure 7. It can be seen that, like other neural networks, these images look like digits.

12 Layer by Layer Classification Results

We also verify the classification results with multiple layer features learned by our model and compare them with the results obtained only by the top layer features. The results are reported in Tables 3 (for MNIST) and 4 (for Caltech 101).

Table 3: Classification Error of MNIST data

Methods	Test error
layer by layer deconvolution	0.42%
One layer deconvolution	0.42%

Table 4: Classification Results (accuracy %) of Caltech 101 data

Training images per class	30
layer by layer deconvolution + all layers feature	84.4±0.1%
One layer deconvolution + one layer feature	84.3±0.3%

13 More Results

Fig. 5 plots dictionaries trained per class and four classes are shown. It can be seen the layer one dictionary elements are pretty similar while the upper layer dictionaries vary from each other. We also performed the model on a joint dataset with multi-class images, with trained dictionary elements shown in Fig. 9. It can be seen that the first layer and second layer dictionaries are more general in this joint training process. For instance, the first layer filters in Fig. 9 is an integration of filters of each class in Fig.5.

Table 5 summarizes the classification accuracy of our model and some related models with 15, 30 and 60 training images per class. So far, only deep convolutional models have provided accuracy greater than 60%. The best result not using deep convolutional models (using hand-crafted features) is 58% by Bo et al., CVPR 2013 (Bo, L., Ren, X., and Fox, D. Multipath sparse coding using hierarchical matching pursuit. CVPR 2013.). Note that Zeiler & Fergus, ECCV2014 used a large dataset, namely, the ImageNet, to train the convolutional network, thus to get good results, while we only use the images inside Caltech 256; this same issue exists for the Zeiler & Fergus, ECCV2014 Caltech 101 results.

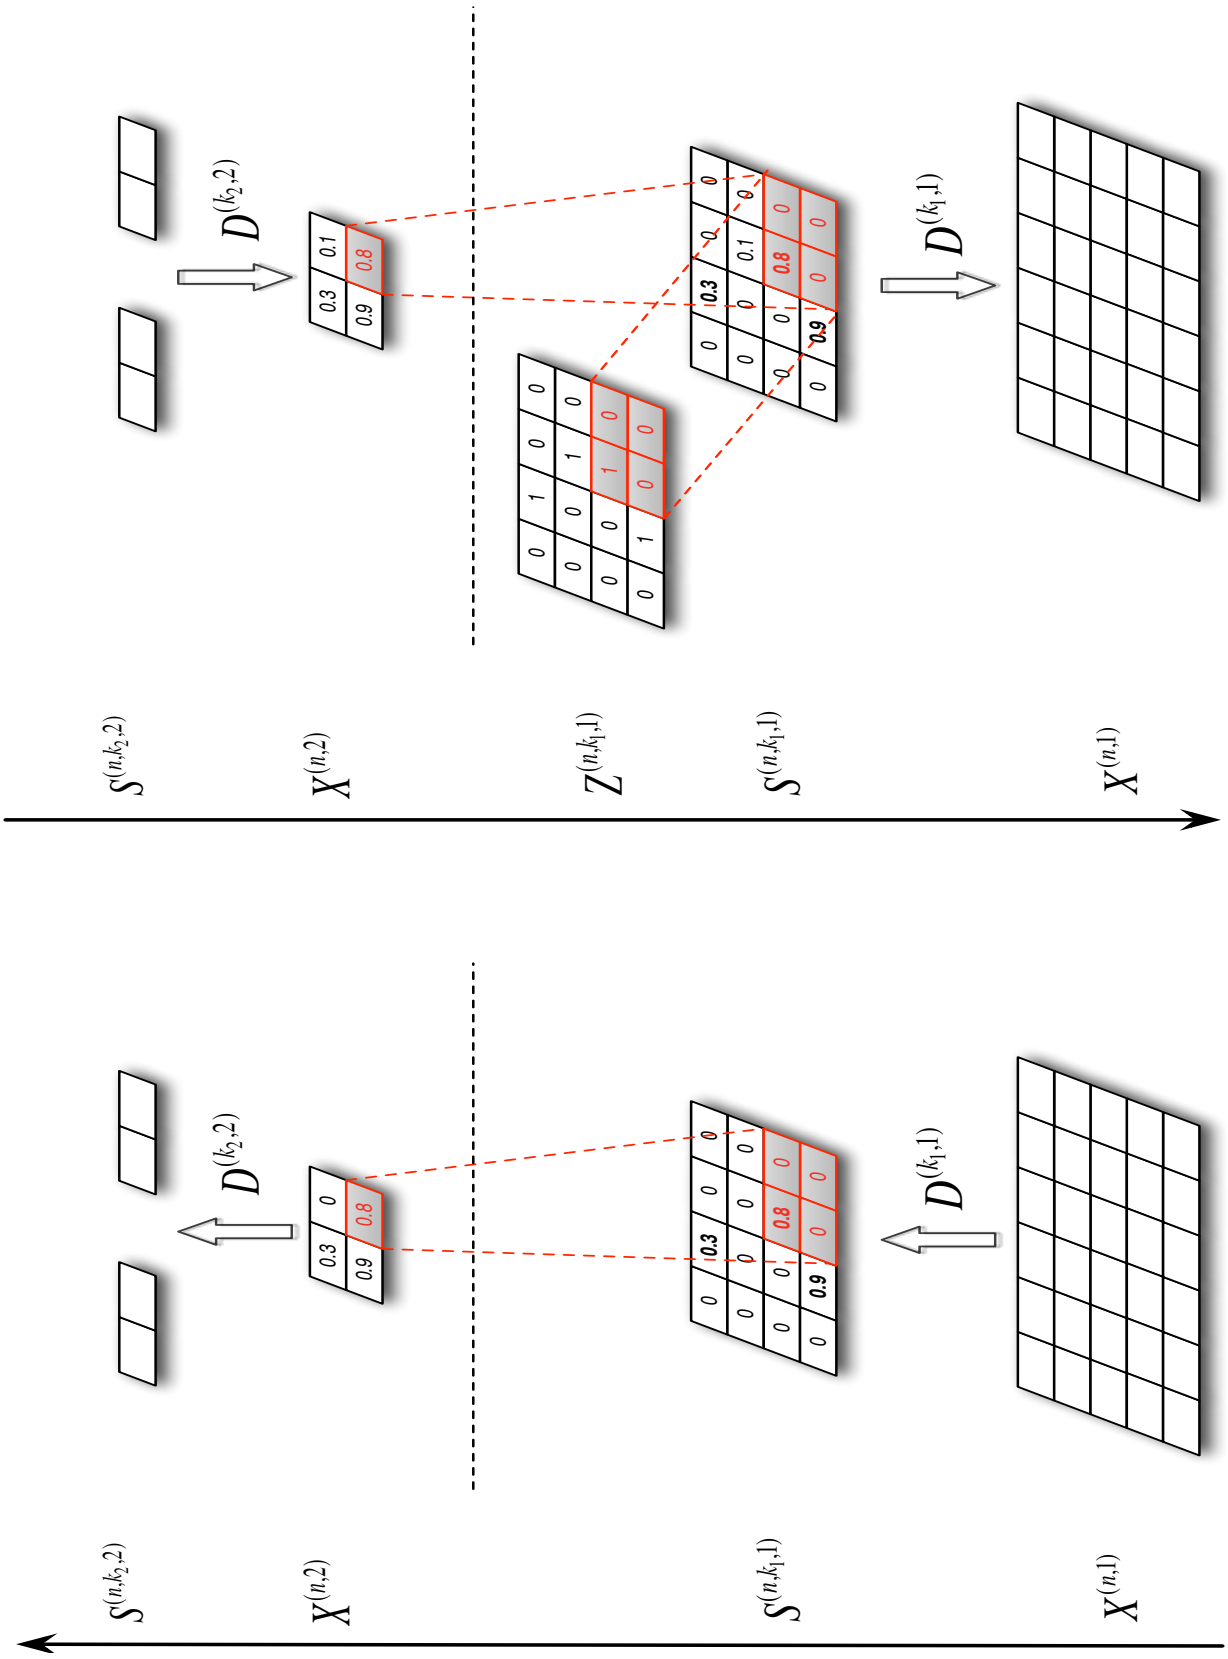

Figure 6: Schematic of the proposed generative process. Left: bottom-up pretraining, right: top-down refinement.

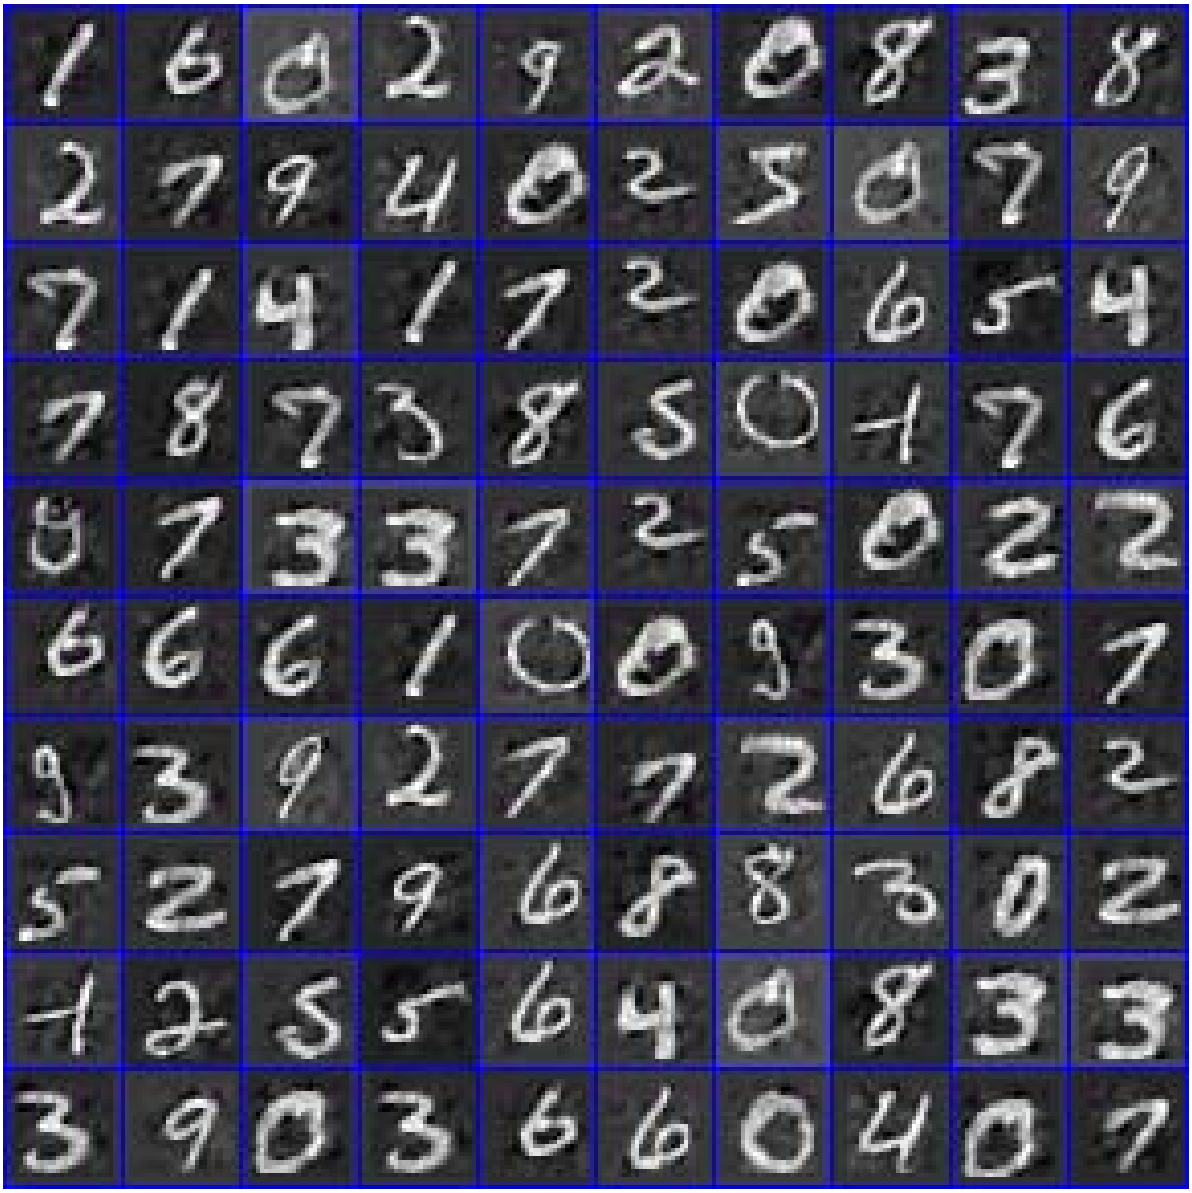

Figure 7: Generated images by randomly drawn dictionary weigths.

Table 5: Classification Accuracy (%) for Caltech 256

Training images per class	**15**	**30**	**60**
Bo et al., CVPR 2013 (Multipath sparse coding)	42.7	48.0	58
Our model	52.4±0.3	58.3±0.2	65.2±0.3
Zeiler & Fergus, ECCV, 2014	65±0.2	70.2±0.2	74.2±0.3

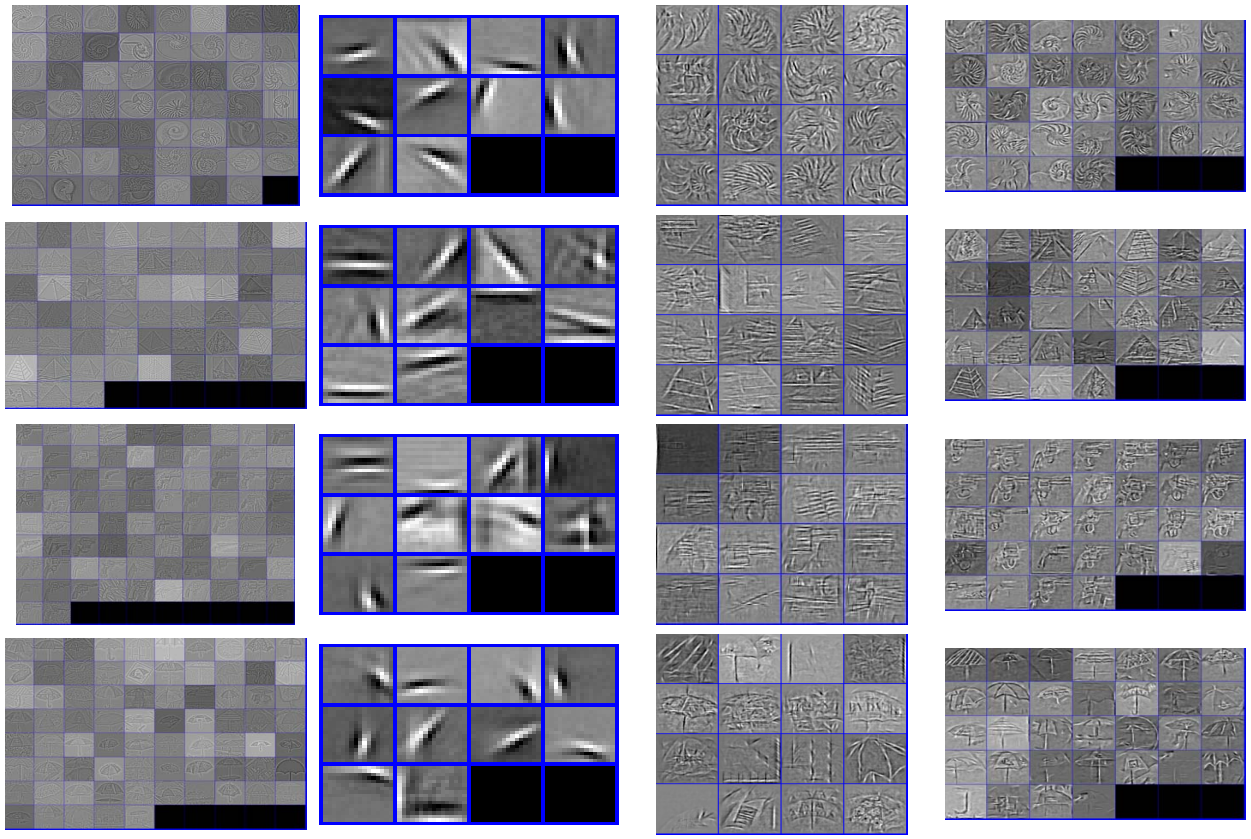

Figure 8: Trained dictionaries per class with each class (from Caltech 101) trained independently. Row 1-4: nautilus, pyramid, revolver, umbrella. Column 1-4: training images after local contrast normalization, layer 1 dictionary, layer 2 dictionary, layer 3 dictionary.

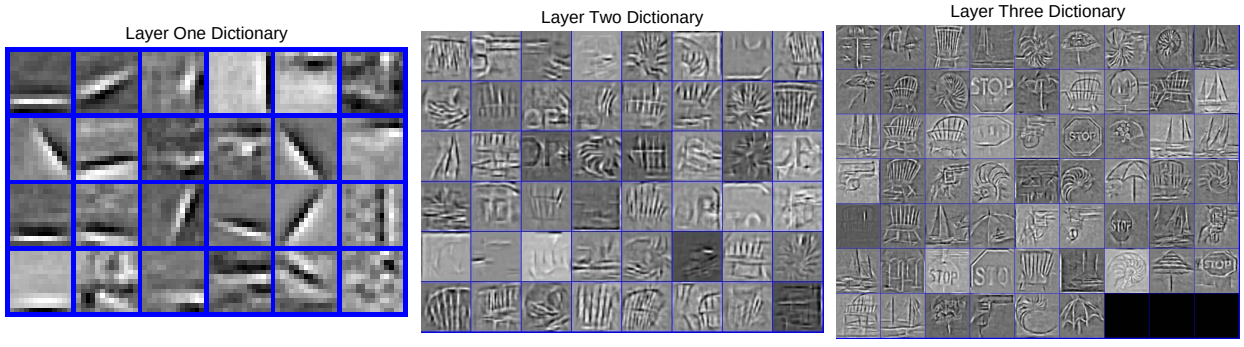

Figure 9: Dictionaries (layer 1 to layer 3) trained by 6 classes (ketch, nautilus, revolver, stop sign, umbrella, windsor chair) images from Caltech 101.

14 DETAILED PARAMETERS USED IN EXPERIMENTS

For MNIST, a 2-layer model ($K_1 = 32$, $K_2 = 160$) is used. For the first layer, we have 32 dictionary nodes (each with size 8×8) and 32×60000 weights (each with size 21×21). For the second layer, we have 160 dictionary nodes (each with size $6 \times 6 \times 32$) and 160×60000 weights (each with size 2×2). These parameters and the latent parameters in the model are list in Table 6.

Table 6: Parameter number used in MNIST Training

Parameter	Number
Layer 1, $\mathbf{D}$	$32 \times 8 \times 8$
Layer 1, $\mathbf{W}$	$32 \times 21 \times 21 \times 60000$
Layer 1, $\mathbf{Z}$	$32 \times 21 \times 21 \times 60000$
Layer 2, $\mathbf{D}$	$160 \times 6 \times 6 \times 32$
Layer 2, $\mathbf{W}$	$160 \times 2 \times 2 \times 60000$
Layer 2, $\mathbf{Z}$	$160 \times 2 \times 2 \times 60000$
Latent Parameter	Number
Layer 1, γ_e	60000
Layer 1, γ_d	32
Layer 1, γ_w	$32 \times 21 \times 21 \times 60000$
Layer 1, $\boldsymbol{\theta}$	$32 \times 21 \times 21 \times 60000$
Layer 2, γ_e	60000
Layer 2, γ_d	160×32
Layer 2, γ_w	$160 \times 2 \times 2 \times 60000$
Layer 2, $\boldsymbol{\theta}$	$160 \times 2 \times 2 \times 60000$

For Caltech 101, in the pretraining step, a 3-layer model ($K_1 = 8$, $K_2 = 12$, $K_3 = 16$) is used for each category. For the first layer, we have 8 dictionary nodes (each with size 17×17) and 8×30 weights (each with size 112×112). For the second layer, we have 12 dictionary nodes (each with size $9 \times 9 \times 8$) and 12×30 weights (each with size 20×20). For the third layer, we have 16 dictionary nodes (each with size $6 \times 6 \times 12$) and 16×30 weights (each with size 5×5). These parameters and the latent parameters in the model are list in Table 7. In the testing, a single-layer model ($K = 1616$) is used. We have 1616×6114 weights (each with size 5×5).

Table 7: Parameter number used in Caltech101 Pretraining(Each Category)

Parameter	Number
Layer 1, $\mathbf{D}$	$8 \times 17 \times 17$
Layer 1, $\mathbf{W}$	$8 \times 112 \times 112 \times 30$
Layer 1, $\mathbf{Z}$	$8 \times 112 \times 112 \times 30$
Layer 2, $\mathbf{D}$	$12 \times 9 \times 9 \times 8$
Layer 2, $\mathbf{W}$	$12 \times 20 \times 20 \times 30$
Layer 2, $\mathbf{Z}$	$12 \times 20 \times 20 \times 30$
Layer 3, $\mathbf{D}$	$16 \times 6 \times 6 \times 12$
Layer 3, $\mathbf{W}$	$16 \times 5 \times 5 \times 30$
Layer 3, $\mathbf{Z}$	$16 \times 5 \times 5 \times 30$
Latent Parameter	Number
Layer 1, γ_e	30
Layer 1, γ_d	8
Layer 1, γ_w	$8 \times 112 \times 112 \times 30$
Layer 1, $\boldsymbol{\theta}$	$8 \times 112 \times 112 \times 30$
Layer 2, γ_e	30
Layer 2, γ_d	12×8
Layer 2, γ_w	$12 \times 20 \times 20 \times 30$
Layer 2, $\boldsymbol{\theta}$	$12 \times 20 \times 20 \times 30$
Layer 3, γ_e	30
Layer 3, γ_d	16×12
Layer 3, γ_w	$16 \times 5 \times 5 \times 30$
Layer 3, $\boldsymbol{\theta}$	$16 \times 5 \times 5 \times 30$